

This Page Is Inserted by IFW Operations
and is not a part of the Official Record

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

**As rescanning documents *will not* correct images,
please do not report the images to the
Image Problem Mailbox.**

LABEL SELECTION FOR END-TO-END LABEL-SWITCHEDTRAFFIC THROUGH A COMMUNICATIONS NETWORKCROSS-REFERENCE TO RELATED APPLICATIONS

This is the first application filed for the present
5 invention.

MICROFICHE APPENDIX

Not Applicable.

TECHNICAL FIELD

The present invention relates to mapping
10 communications paths through a communications network, and
in particular to label selection for end-to-end label
switched traffic through a communications network.

BACKGROUND OF THE INVENTION

Optical communications networks have recently
15 become established as the preferred backbone for data
communications because of the high bandwidth capacity of
the optical fiber links. In the modern network space,
packetized data traffic of multiple different protocols
(e.g. internet protocol, frame relay, asynchronous transfer
20 mode, etc.) is transported over a common network
infrastructure. Each protocol provides its own packet (or
frame) size and format standards. Additionally, some
protocols (e.g. IP) are specifically designed to allow
packets having widely varying lengths. New routing
25 protocols, for example the multi-protocol label switching
(MPLS) protocol have been proposed to facilitate
multi-protocol traffic across a common network
infrastructure.

Under the MPLS protocol, label switched packets (LSPs) are propagated across the network hop-by-hop along a path that is set up at the beginning of a communications session. In general, the label assigned to the LSP can be
5 different for each hop, with the label conversion being performed by the node serving the respective hop. Where the network is designed for wave division multiplex (WDM) transport of data traffic, it is desirable to use a data channel (or wavelength) as the label assign to an LSP
10 through a respective hop. In this case, end-to-end transport of an LSP through an MPLS path requires a change of wavelength in the node serving each hop.

However, the modern network space is composed of a mixture of agile cross-connects (i.e. cross-connects
15 capable of wavelength conversion from input to output) and non agile cross-connects (i.e. cross-connects that are unable to do wavelength conversion). Because MPLS normally requires that the labels (i.e. the wavelengths) must be changeable at every hop, non agile cross-connects within
20 the network cannot participate in an MPLS end-to-end optical path. This restriction forces MPLS traffic to be routed around non-agile cross-connects, and serves as a barrier to the widespread implementation of MPLS.

A technique which allows non-agile cross-connects
25 to participate in MPLS end-to-end optical paths is therefore highly desirable.

SUMMARY OF THE INVENTION

An object of the present invention is to provide a technique which enables an MPLS end-to-end optical path to
30 be mapped through a communications network, in which at

least one and possibly all of the nodes participating in the path are non-agile.

A further object of the present invention is to provide a technique enabling an MPLS end-to-end optical
5 path to be mapped across a communications network using a common label (e.g. wavelength) for every hop of the path, so that inter-hop label conversion is not necessary.

Accordingly, an aspect of the present invention provides a method of label selection for end-to-end
10 transport of label switched traffic through a communications network between a source node and a destination node. A request message is launched toward the destination node from the source node. The request message includes a label list having one or more label identifiers
15 indicative of respective corresponding labels available for use by the source node. The label list is revised, at a cross-connect service each successive hop between the source node and the destination node, based on labels available for use by each respective hop, to produce a
20 reduced label list. The reduced label list includes label identifiers indicative of respective corresponding labels available for end-to-end transport of label switched traffic between the source node and the destination node.

A further aspect of the present invention provides
25 communications network adapted for end-to-end transport of label switched traffic through the communications network. The communications network comprises a source node, a destination node, and at least one cross-connect intermediate the source node and the destination node. The
30 source node is adapted to launch a request message toward a destination node. The request message includes a label list having one or more label identifiers indicative of

respective corresponding labels available for use by the source node. Each cross-connect is adapted to revise the label list based on labels available for use over a successive hop served by the cross-connect, to produce a
5 reduced label list.

A still further aspect of the invention provides a cross-connect of a communications network adapted for end-to-end transport of label switched traffic through the communications network between a source node and a
10 destination node. The cross-connect includes a label availability table, a buffer and a processor. The label availability table includes label identifiers indicative of labels available for a successive hop served by the cross-connect. The buffer is adapted to receive a request
15 message propagated through the communications network from the source node. The request message includes a label list having label identifiers indicative of labels available for conveying label switched traffic between the source node and the cross-connect. The processor is adapted to revise
20 the label list included in the received request message, based on the label availability table, to produce a reduced label list.

The communications network may be adapted for wave division multiplex (WDM) transport of label switched
25 traffic, each label comprising a respective data transport wavelength.

The label switched traffic may include multi-protocol label switched (MPLS) traffic.

In an embodiment of the invention, the step of
30 revising the label list comprises, at each cross-connect serving a successive hop toward the destination node, a

step of intersecting the label list with a set of label identifiers indicative of labels available for use over the hop. A request rejection message may be sent to the source node if the reduced label list is empty.

5 If the reduced label list included with the request message received by the destination node contains at least one label identifier, then one of the label identifiers is selected from the reduced label list. An end-to-end label switched path is then set up between the source node and
10 the destination node using the respective label corresponding to the selected label identifier. The label identifier may be selected at random from the reduced label list.

 The step of setting up an end-to-end label switched
15 path may include sending a mapping message containing the selected label identifier from the destination node toward the source node, the mapping message retracing the path traversed by the request message. Upon receipt of the mapping message at the cross-connect serving each hop, the
20 corresponding label indicated by the selected label identifier is assigned to the end-to-end label switched path, if the label is still available for use by the hop.

 A mapping failure message may be sent to the destination node if the label corresponding to the selected
25 label identifier is not available for use by the hop. In this case, upon receipt of the mapping failure message by the destination node, the reduced label list may be stripped by removing the selected label identifier. If the thereby stripped label list is empty, a request rejection
30 message may then be sent to the source node. Otherwise, if the stripped label list contains at least one label identifier, a new label identifier may be selected from the

stripped label list, and an end-to-end label switched path set up between the source node and the destination node using the respective label corresponding to the newly selected label identifier.

5 **BRIEF DESCRIPTION OF THE DRAWINGS**

Further features and advantages of the present invention will become apparent from the following detailed description, taken in combination with the appended drawings, in which:

10 Fig. 1 is a block diagram illustrating an optical communications network usable in conjunction with an embodiment of the present invention;

Fig. 2 is a block diagram schematically illustrating principal elements of a cross-connect usable
15 in the optical communications network illustrated in Fig. 1; and

Fig. 3 is a message flow diagram schematically illustrating principal messages exchanged during the set up of an end-to-end optical path in accordance with an
20 embodiment of the present invention.

It will be noted that throughout the appended drawings, like features are identified by like reference numerals.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

25 As shown in Fig. 1, an optical network 2 usable in conjunction with the present invention generally comprises a plurality of cross-connects 4 (four are shown in Fig. 1) interconnected by fiber optic links 6. Communications devices 8, for example end user personal computers (PCs) or

local area network (LAN) servers may be connected to the optical network 2 via one or more edge nodes (or access points) 10. The optical network 2 may also be connected to one or more federated networks 12, for example an asynchronous transfer mode (ATM) or an internet protocol (IP) network, through a respective gateway 14. Within the optical network 2, each of the cross-connects 4 is configured for wave division multiplex (WDM) and/or dense wave division multiplex (DWDM) transport of packet data traffic as will be described in greater detail below.

The present invention provides a technique for label selection for an MPLS end-to-end optical path 16 across the communications network 2 between a source node 10a and an destination node 10b via one or more intervening cross connects 4. The path 16 is divided into hops 18, each of which is served a respective node (e.g. the source node 10a of a cross-connect 4) connected at the up-stream end of the hop 18. In a WDM (or DWDM) environment, the label assigned to each label switched packet (LSP), for each hop 18, is the channel (wavelength) used for conveying the LSP through the hop. In other words, the label designates the wavelength on which the path 16 traverses the hop 18. In the example illustrated in Figs. 1 and 3, the source and destination nodes 10a and 10b are located at respective edge nodes, and two intervening cross-connects 1a and 4b are incorporated into the path 16. One or both of the cross connects 4a,b may be non-agile, so that label (thus wavelength) conversion between hops (i.e. within a cross-connect 4) may not be possible. Thus in accordance with the present invention, a common label is selected such that the label switched end-to-end optical path 16 using the common label

(wavelength) can be mapped across the communications network 2.

Fig. 2 is a block diagram schematically illustrating the principal elements of a cross-connect 4 of the optical network. As shown in Fig. 2, the cross-connect 4 includes at least one input port 20; at least one output port 22; a switch 24 (which may be an optical space switch) capable of providing a signal path between respective input and output ports 20, 22; a controller 26 for controlling operation of the cross-connect 4; and a label availability table 28, which may be co-resident with the cross-connect 4 or located at a remote site and accessible by the controller 26. The controller 26 can be provided with a buffer (not shown) for temporarily storing information.

In the embodiment of Fig. 2 the cross-connect 4 is provided with two input ports 20a and 20b and two output ports 22a and 22b. However, it will be appreciated that a greater or smaller number of ports may be provided in the cross-connect 4, and/or the cross-connect 4 may be provided with a plurality of ports of which one or more may remain unused. It will also be appreciated that the ports of the cross-connect 4 may be suitably designed and connected to handle bi-directional data traffic. However, in order to simplify the present description, and aid understanding of the invention, the embodiment of Fig. 2 is provided with unidirectional ports, two of which (input ports 20a and 20b) are connected to handle inbound LSPs, and the other two (output ports 22a and 22b) are connected to handle outbound LSPs.

The label availability table 28 dynamically stores information concerning the availability of each label which

can be used for transporting LSPs over hops 18 immediately downstream of the cross-connect 4. Thus for each output port 22, the label availability table 28 contains a list of label identifiers corresponding to each label (i.e. transmission wavelength), and an associated availability flag. The availability flag may conveniently be a binary value, in which a "1" indicates that the respective label indicated by the corresponding label identifier is available for use, and a "0" indicates that the respective label is not available for use. The label availability table 28 can be dynamically maintained by updating the availability flag as label switched paths 16 are set up and released. In the embodiment illustrated in Figs. 2 and 3, four labels (indicated by respective label identifiers λ_1 through λ_4) may be used for label switched traffic. Clearly, more, or fewer labels may be used, as desired.

Operation of the cross connect 4, and the method of the present invention, is described in greater detail below with reference to Fig. 3.

As shown in Fig. 3, the source node 10a is provided with a respective label availability table 30 which is closely similar to that of each of the cross-connects 4a and 4b involved in the path 16. In the example illustrated in Fig. 3, the label availability table 30 of the source node 10a indicates that all possible labels (λ_1 through λ_4) are available for use on the first hop 18a of the path 16 (i.e. between the source node 10b and the first cross-connect 4a). The label availability table 28a associated with the first cross-connect 4a indicates that, of the possible labels λ_1 through λ_4 , only the first three labels (λ_1 through λ_3) are available for use on the second

hop 18b of the path 16 (i.e. between the first and second cross-connects 4a and 4b). Finally, the label availability table 28b associated with the second cross-connect 4b indicates that, of the possible labels λ_1 through λ_4 , only
5 the first and third labels (i.e. λ_1 and λ_3) are available for use on the third hop 18c of the path 16 (i.e. between the second cross-connect 4b and the destination node 10b).

In order to select a label and set up the end-to-end path 16, the source node 10a prepares a request
10 message R (step 100) which includes a label list 32 identifying each of the labels which are available for use by the source node 10a for transporting LSPs over the first hop 18a of the path 16. Conveniently, the label list 32 can be prepared by extracting each label identifier from the
15 label availability table 30 for which the associated availability flag is set to "1". The request message R including the label list 32 is then launched over a link 6 (step 102) towards the first cross-connect 4a en-route to the destination node 10b.

20 Upon receipt of the request message R, the first cross-connect 4a extracts the label list 32 and compares it with its associated label availability table 28a (step 104). The first cross-connect 4b performs this comparison operation by calculating the intersection
25 between the set of labels identified in the label list 32 and the set of labels available for the next hop 18b of the path 16, as indicated by the label availability table 28a. The intersection set resulting from this analysis is the set of label identifiers corresponding to labels which are
30 available for use over the first two hops 18a,b of the path 16. This intersection set is then inserted into the request message R as a reduced label list 34, and the

request message R forwarded on to the second cross-connect 4b (step 106).

Upon receipt of the request message R including the reduced label list 34, the second cross-connect 4b extracts
5 the reduced label list 34 and compares it with its associated label availability table 28b (step 108) in an identical manner to that described above with respect to the first cross-connect 4a. The resulting intersection set includes each label identifier corresponding to labels
10 which are available for use over all three hops 18a-18c of the path 16. This intersection set is inserted into the request message R as a further reduced label list 34b, and forwarded to the destination node 10b (step 110).

It will be appreciated that this process of
15 receiving the request message R; extracting the label list, and calculating the intersection of the label list and the set of available label identifiers in an associated label availability table 28 can be repeated (over any number of hops 18) until the reduced label list is empty or the
20 request message (containing a non-empty label list) is forwarded to the destination node 10b. At any cross-connect 4, if the reduced label list is found to be empty, then a request rejection message (not shown) can be generated and sent back to the source node 10a to indicate
25 that no labels are available for the end-to-end path 16).

Returning now to Fig. 3, upon receipt of the request message R, the destination node 10b extracts the reduced label list 34b, and determines whether or not it contains any label identifiers (step 112). In the present
30 example, the reduced label list 34b contains two label identifiers, namely λ_1 and λ_3 . The destination node 10b

then selects one of the label identifiers (e.g. by random selection), and launches a mapping message M (step 114) including the selected label identifier λ_1 back to the source node 10a, and retracing the route followed by the request message R.

As the mapping message M propagates hop-by-hop towards the source node 10a, the label corresponding to the selected label identifier λ_1 is assigned by each cross-connect 4 to the end-to-end path 16 and the corresponding entry in the respective label availability table 28 is updated. Thus in the example illustrated in Fig. 3, the mapping message M including the selected label identifier λ_1 is forwarded by the destination node 10b over the third hop 18c. Upon receipt of the mapping message M, the second cross-connect 4b extracts the selected label identifier λ_1 ; assigns the corresponding label to the end-to-end path 16; and updates its label availability table 28b to indicate that the label is no longer available (step 116). The mapping message M including the selected label identifier λ_1 is then forwarded over the second hop 18b to the first cross-connect 4a (step 118).

It will be appreciated that it is possible that in the period of time between reception of the request message R propagated from the source node 10a, and subsequent reception of the mapping message M propagated from the destination node 10b, a cross-connect 4 may have received another mapping message in respect of a different path and including the same label identifier. In this case, the label corresponding to the selected label identifier λ_1 will no longer be available, upon subsequent arrival of the mapping message M propagated from the destination node 10b.

This situation is illustrated in Fig. 3. Thus upon receipt of the mapping message M, the first cross-connect 4a extracts the selected label identifier λ_1 , and attempts to assign the corresponding label to the path 16. However, the
5 corresponding label has already been assigned to another path, and is therefore no-longer available. In this circumstance, the first cross-connect 4a prepares a mapping failure message MF including the selected label identifier λ_1 and indicating that the corresponding label is no longer
10 available for use. The mapping failure message MF is then forwarded (steps 122-126) hop-by-hop back to the destination node 10b, so that resources already assigned to the path 16 en-route to the destination node 10b can be released.

15 Upon receipt of the mapping failure message MF including the selected label identifier λ_1 , the destination node 10b extracts the label identifier λ_1 and strips it from the reduced label list 34b (step 128) previously received with the request message R. If the thereby
20 stripped label list 34c is empty, then the destination node 10b forwards a request rejection message (not shown) back to the source node 10a indicating that there are no labels available for the end-to-end path 16. Otherwise, the destination node 10b selects (e.g. at random) another label
25 identifier from the stripped label list 34c (step 130), and attempts to set up the end-to-end path 16 by launching a new mapping message M' including the newly selected label identifier toward the source node 10a and retracing the route followed by the request message R. In the present
30 example, the stripped label list contains a single label identifier (λ_3), and this label identifier is inserted into the new mapping message M' which is launched towards the

source node 10a (steps 132, 136 and 140). At each cross-connect 4, resources are assigned to the path 16 (steps 134 and 138) and the associated label availability table 28 updated. Arrival of the mapping message at the
5 source node 10a (at step 142) including the newly selected label identifier λ_3 indicates completion of the end-to-end path 16 through the communications network 2.

Thus the present invention provides a simple means by which an end-to-end label switch path having a common
10 label (in the present example, the label indicated by label identifier λ_3) for each hop of the path is set up. Consequently, non-agile cross-connects can participate in the end-to-end path, because it is not necessary to perform wavelength conversions between hops.

15 The embodiment(s) of the invention described above is(are) intended to be exemplary only. The scope of the invention is therefore intended to be limited solely by the scope of the appended claims.